arXiv:1702.00171v1 [cs.IR] 1 Feb 2017

Inferring Conceptual Relationships when Ranking Patients

Nut Limsopatham[a], Craig Macdonald[b], Iadh Ounis[b]

[a]*Language Technology Lab, University of Cambridge, United Kingdom*
[b]*School of Computing Science, University of Glasgow, United Kingdom*

Abstract

Searching patients based on the relevance of their medical records is challenging because of the inherent implicit knowledge within the patients' medical records and queries. Such knowledge is known to the medical practitioners but may be hidden from a search system. For example, when searching for the patients with a heart disease, medical practitioners commonly know that patients who are taking the amiodarone medicine are relevant, since this drug is used to combat heart disease. In this article, we argue that leveraging such implicit knowledge improves the retrieval effectiveness, since it provides new evidence to infer the relevance of patients' medical records towards a query. Specifically, built upon existing conceptual representation for both medical records and queries, we proposed a novel expansion of queries that infers additional conceptual relationships from domain-specific resources as well as by extracting informative concepts from the top-ranked patients' medical records. We evaluate the retrieval effectiveness of our proposed approach in the context of the TREC 2011 and 2012 Medical Records track. Our results show the effectiveness of our approach to model the implicit knowledge in patient search, whereby the retrieval performance is significantly improved over both an effective conceptual representation baseline and an existing semantic query expansion baseline. In addition, we provide an analysis of the types of queries that the proposed approach is likely to be effective.

1. Introduction

Electronic medical records (EMRs) detail the medical history of patients visiting healthcare providers [1, 2]. Within a patient search system, these EMRs could be exploited to identify cohorts for clinical studies [3, 4, 5]. One of the major challenges of searching in the medical domain is to deal with often complex, inconsistent and ambiguous terminology [6, 7, 8]. For example, it is commonly known by medical practitioners that *'cancer'*, *'carcinoma'*, *'CA'*, and *'malignant tumour'* share a similar meaning. However, such information may be hidden from traditional search systems. To handle such a challenge, prior works resorted to domain-specific resources to improve the representation of medical documents and queries. For instance, by exploiting the knowledge obtained from domain-specific resources, such as MeSH,[1] using concept-based representation approaches, *'cancer'*, *'carcinoma'*, *'CA'*, and *'malignant tumour'* are represented with the same concept to lessen the mismatch of synonymous terms in medical documents and

Email addresses: nl347@cam.ac.uk (Nut Limsopatham), craig.macdonald@glasgow.ac.uk (Craig Macdonald), iadh.ounis@glasgow.ac.uk (Iadh Ounis)

[1] http://www.nlm.nih.gov/mesh/

queries [8, 9, 10]. Moreover, synonym and hyponym relationships of medical terms obtained from domain-specific resources have been effectively exploited to reformulate queries in order to further improve retrieval performance (e.g. [11, 12]). Hence, a search system could infer that patients whose medical records stated that they are suffering from *'pulmonary atresia'* are relevant to a query searching for the patients suffering from *'a heart disease'*, since according to MeSH, *'pulmonary atresia'* is a particular form of *'a heart disease'*.

Nevertheless, the aforementioned approaches do not leverage the inherent implicit knowledge in medical records and queries, which could be exploited to improve retrieval effectiveness. For example, we can infer that a patient is suffering from a particular disease, if the patient is prescribed a medicine used specifically to treat that disease. Statistical query expansion (QE) approaches, such as pseudo-relevance feedback, indirectly deal with the implicit knowledge challenge by using occurrence statistics of terms in the top-ranked documents to improve the representation of the original query [13]. For example, Limsopatham et al. [14, 15] effectively applied the Divergence from Randomness (DFR) Bo1 QE model [13] to improve the retrieval performance of a patient search system. However, these QE approaches might not effectively deal with the implicit knowledge challenge, if inherently related concepts are not observed in the top-ranked records.

Instead, to leverage the implicit knowledge inherent to the patient search process, in this article, we go beyond classical statistical QE techniques. We propose to improve the representation of medical queries by inferring the relationships of medical conditions that are extracted from the queries. In particular, we focus on medical conditions that are related to *four aspects of the medical decision process (namely, symptom, diagnostic test, diagnosis, and treatment)* [16, 3, 17], as they are important information considered by healthcare practitioners when consulting with a patient. Next, using this conceptual representation of the medical records and queries, we propose a novel query expansion approach to further improve the query representation by exploiting two types of resources. The first type consists of external resources about the possible relationships between medical conditions (i.e. medical concepts). For example, we can infer that a patient suffers from a particular disease, if the patient is prescribed a medicine for that disease. To infer the relationships between medical concepts using external resources, we propose two techniques, which are based on Bayes' theorem and a stochastic analysis, respectively. On the other hand, as a second type of resource, we use the collection of patients' medical records itself. In particular, we extract the most informative concepts from the top-ranked medical records retrieved using the original query. The medical concepts inferred using these two types of resources are used to expand the original queries to improve their representations.

We evaluate our proposed approach in the context of the TREC 2011 [4] and 2012 [5] Medical Records track. Our results attest the effectiveness of our proposed approach, as it can significantly improve the retrieval performance over an existing query expansion baseline.

The main contributions of this paper are threefold:

1. We propose a novel query expansion approach that models the relationships between concepts using the four aspects of medical decision process, by leveraging medical knowledge gained from external resources and the occurrence statistics of concepts in the top-ranked medical records to improve the query representation and to infer relevance.

2. We introduce two techniques for inferring the relationships between medical concepts derived from external resources, based on Bayes' theorem and a stochastic analysis, respectively.

3. We thoroughly evaluate our proposed approach within the standard experimentation paradigm provided by the TREC 2011 and 2012 Medical Records track.

The remainder of this paper is organised as follows. In Section 2, the backgrounds of searching patients based on the relevance of their medical records and related works are discussed. Section 3 discusses our approach for improving query representation that infers relationships between medical concepts by using information from both external resources and the informative concepts extracted from the top-ranked medical records. Experimental setup and results are presented in Sections 4 and 5, respectively. Section 6 provides a further analysis and discussion of the proposed approach. Finally, we conclude the paper in Section 7.

2. Related Work

Electronic medical records (EMRs), which detail the healthcare information of patients, have been developed to improve the quality of healthcare services [18]. For instance, EMRs could be exploited to identify treatments that have been used effectively to combat a particular disease [18, 19]. However, the characteristics of medical records and queries, such as the complexity of the medical terminology, are different from those of other domains. Hence, effective search approaches for medical records are needed. In 2011, TREC developed a search task to facilitate the research in this area [4]. In particular, the task of the TREC Medical Records track [4, 5] aims to rank patients with respect to the relevance of their medical records towards a query.

Prior work (e.g. [20, 21, 22]) effectively handled this search task using techniques previously developed for expert search [23], since the goal of both tasks is to rank people (i.e. patients or expert persons) based on the relevance of their associated documents. On one hand, expert search aims to rank experts based on the relevance of the documents they have written or that mention them [23]. On the other hand, patient search uses the estimated relevance of medical records, when ranking their associated patients. Hence, in this work, we also handle patient search using well-established approaches previously developed for expert search, which use ranked medical records to rank patients (e.g. Voting Model [24] and Model 2 [25]). Specifically, the Voting Model ranks patients using a voting process, where the ranking of medical records (denoted $R(Q)$) defines the relevance scores for the patients to be retrieved. Each retrieved medical records in $R(Q)$ is said to vote for the relevance of its associated candidate patient using voting techniques such as, CombMAX, CombSUM, expCombSUM. Indeed, each voting technique firstly ranks medical records based on their relevance towards a query using a traditional weighting model (e.g. BM25 [26], DFR DPH [27]), and then aggregates the votes from the medical records to their associated patients, to create a ranked list of likely relevant patients for the query [24].

One of the important research areas of searching in the medical domain is dealing with the complexity, ambiguity and inconsistency of the medical terminology [6, 7, 8]. For example, when referring to *'coronary heart disease'*, different medical practitioners may use terms, such as *'coronary artery disease'*, *'arteriosclerotic heart disease'*, *'CHD'*, or *'CAD'*. Previous work has leveraged domain-specific resources to handle such a challenge [6, 7, 8]. For instance, Srinivasan [10] and Trieschnigg et al. [8] represented medical documents and queries using medical concepts obtained from domain-specific resources, such as MeSH, to alleviate the synonymous mismatch of terms in medical documents and queries. Aronson [28] deployed MetaMap [29] – a medical concept recognition tool based on the UMLS Metathesaurus[2] – to identify all concept,

[2] http://www.nlm.nih.gov/research/umls/

in the medical documents and queries, and to represent them in the form of the UMLS Concept Unique Identifier (CUI). However, conceptual representation approaches are effective only when combined with a traditional term-based representation [8, 10]. Koopman et al. [30] proposed a graph-based approach for weighting patients based on the occurrences of medical concepts in their medical records. They built a graph model based on the co-occurrences of medical concepts within a particular window of concepts, and calculated the PageRank [31] scores for each medical concepts. Then, during retrieval, they scored patients based on the PageRank scores of the query concepts that matched their medical records. Later, Limsopatham et al. [3] proposed a task-specific representation approach, which represented medical records and queries by using only concepts that are related to four aspects of the medical decision process (namely, symptom, diagnostic test, diagnosis and treatment) [16], as they are essential information for health practitioners when dealing with a patient. Note that these aspects can be viewed as the criteria aspects that the medical records of relevant patients should cover as suggested in [32, 33]. In this work, we follow [3] when representing medical records and queries. In particular, we deploy MetaMap to extract concepts, in medical records and queries, and represent the identified concepts in the form of the UMLS Concept Unique Identifier (CUI). Importantly, we identify concepts related to the four aforementioned aspects based on their MetaMap's semantic type.[3] In Table 1, we list the 16 MetaMap's semantic types that are associated to the medical decision process [17]. Medical concepts with the MetaMap's semantic type defined in the first column are associated to an aspect of the medical decision process, if there is a tick (✔) in the column of that aspect. For example, concepts having the MetaMap's semantic type *Disease or Syndrome* are associated with the *diagnosis* aspect. Figure 1 shows an example of using the MetaMap tool to extract medical concepts from the query "Patients with diabetes mellitus who also have thrombocytosis", where we can extract four different CUIs, including 'C0011849', 'C0011847', 'C0241863' and 'C0836924', from. Nevertheless, even though this approach could effectively represent medical concepts mentioned in documents and queries, related medical concepts that are not explicitly mentioned may not be inferred and uncovered.

On the other hand, query expansion (QE) techniques, such as pseudo-relevance feedback, have been used to effectively improve the representation of queries in different search tasks [11, 12, 13]. In the form of QE, synonyms and hyponyms of concepts in the medical documents and queries have been used to improve the representation of medical queries [11, 34, 12, 35]. For example, Aronson and Rindflesch [11], and Srinivasan [12] effectively expanded concepts in a query with their synonyms and hyponyms obtained from domain-specific resources, such as MeSH and UMLS Metathesaurus. In the context of patient search, King et al. [36] effectively deployed the UMLS Metathesaurus and their collection of medical reference encyclopedias in their query expansion technique. In particular, they simply added the derived medical concepts to the query and set the weights equally for all the medical concepts. Zuccon et al. [35] investigated into exploiting hierarchical structure of SNOMED-CT ontology[4] to extract expansion concepts and several techniques (e.g. tf-idf) for calculating weights of the expansion concepts. However, they found a limited performance improvement. Recently, Martinez et al. [34] achieved significant improvement when applied PageRank random walks on the UMLS Metathesaurus to extract expansion concepts and their weight. However, a limitation is that some types of relationships between medical concepts (e.g. the relationships between diagnosis and treatment) may not be directly presented in the UMLS Metathesaurus. Meanwhile, pseudo-relevance feedback is to

[3] `http://metamap.nlm.nih.gov/SemanticTypeMappings_2011AA.txt`

[4] `https://www.nlm.nih.gov/snomed/`

```
Input: "Patients with diabetes mellitus who also
        have thrombocytosis"

Phrase: "Patients"

Phrase: "with diabetes mellitus"
Meta Candidates (4):
  1000 C0011849:Diabetes Mellitus [Disease or Syndrome]
          Diabetes
   861 C0011847:Diabetes [Disease or Syndrome]
   789 C0241863:DIABETIC [Finding]
Meta Mapping (1000):
  1000 C0011849:Diabetes Mellitus [Disease or Syndrome]

Phrase: "who also"

Phrase: "have"

Phrase: "thrombocytosis."
Meta Candidates (1):
  1000 C0836924:Thrombocytosis [Disease or Syndrome]
Meta Mapping (1000):
  1000 C0836924:Thrombocytosis [Disease or Syndrome]
```

Figure 1: Medical concepts extracted by the MetaMap tool from the query 'Patients with diabetes mellitus who also have thrombocytosis'

expand a query with *a set of informative terms* obtained from the top-ranked documents. In the context of searching the medical domain, Srinivasan [12] and Limsopatham et al. [14] reported that pseudo-relevance feedback could effectively improve retrieval performance. In particular, pseudo-relevance feedback can indirectly deal with the implicit knowledge, since it may expand the query with semantically-related concepts. For example, for a query searching for *"patients with heart disease"*, the expanded terms might be *'amiodarone'* and *'angina'*, which are a treatment and a symptom associated to *'heart disease'*. Importantly, in medical records such relationships between concepts related to the aspects of the medical decision process are strongly established. On the other hand, from a medical record of a patient taking *'amiodarone'* medicine (treatment), healthcare practitioners can infer that the patient is suffering from *'heart disease'* (diagnosis), since according to the indication *'amiodarone'* is used to combat *'heart disease'*. However, pseudo-relevance feedback may not always be able to leverage such implicit knowledge, if the associated concepts do not appear in the top-ranked medical records. Hence, in this work, we propose to also directly infer these relationships to improve the query representation. Specifically, we propose a new QE approach that uses conceptual relationships extracted from both external resources and the top-ranked medical records to improve retrieval effectiveness.

Table 1: List of 16 of the MetaMap's 133 semantic types that we consider for our proposed approach, based on the four aspects of the medical decision process.

MetaMap's Semantic Type	Aspects of the Medical Decision Process			
	Symptom	Diagnostic test	Diagnosis	Treatment
Body Location or Region	✔	✔	✔	✔
Body Part, Organ, or Organ Component	✔	✔	✔	✔
Clinical Drug	–	–	–	✔
Diagnostic Procedure	–	✔	–	–
Disease or Syndrome	–	–	✔	–
Finding	✔	–	–	–
Health Care Activity	–	✔	–	✔
Injury or Poisoning	✔	–	–	–
Intellectual Product	–	✔	–	✔
Medical Device	–	✔	–	✔
Mental or Behavioral Dysfunction	✔	–	✔	–
Neoplastic Process	✔	✔	✔	✔
Pathologic Function	✔	–	–	–
Pharmacologic Substance	–	–	–	✔
Sign or Symptom	✔	–	–	–
Therapeutic or Preventive Procedure	–	–	–	✔

3. An Approach for Inferring Conceptual Relationships when Ranking Patients

In this section, we introduce our proposed approach for inferring the relationships between medical concepts that leverage both external resources and the local statistics of the top-ranked medical records. In particular, Section 3.1 discusses how we derive possible relationships between medical concepts from external resources and represent them as association rules. In Section 3.2 and 3.3, we propose two techniques based on Bayes' theorem and a stochastic analysis, respectively, to infer relationships between medical concepts from the extracted association rules. Section 3.4 discusses our approach to improve the representation of the query by using both medical concepts derived from the external resources using either Bayes' theorem or a stochastic analysis, and informative concepts extracted from the top-ranked medical records.

3.1. Extracting and Representing Medical Concept Relationships from External Resources as Association Rules

Driven by the aforementioned medical decision process [17], we extract directed association rules representing the relationships between concepts from two different types of medical resources[5], which are ontology-based and free-text-based resources, respectively. We use different strategies for extracting conceptual relationships from each type of resource. For the ontology-based resources (e.g. MedDRA[6] and DOID[7]), we use the semantic relationships of concepts within each ontology to represent the relationships between concepts. For instance, 'Coronary

[5]We provide the list of all of the used resources in Section 4.

[6]`http://www.meddramsso.com`

[7]`http://purl.bioontology.org/ontology/DOID`

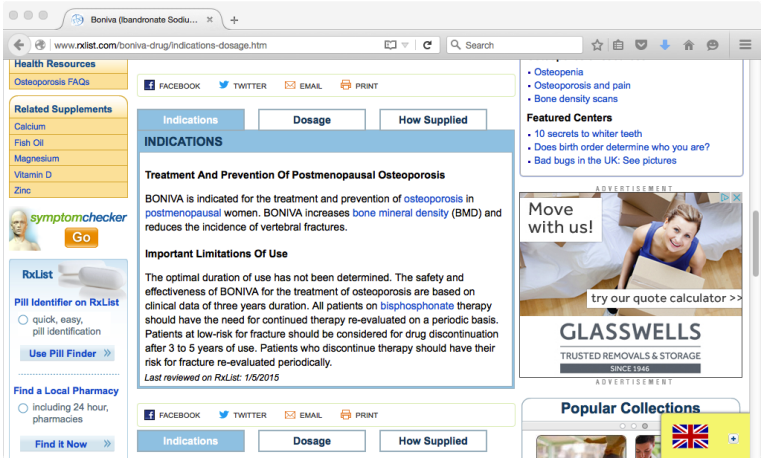
Boniva (Ibandronate Sodiu...
www.rxlist.com/boniva-drug/indications-dosage.htm
Search
Health Resources
Osteoporosis FAQs
Related Supplements
Calcium
Fish Oil
Magnesium
Vitamin D
Zinc
symptomchecker
Go
RxList
Pill Identifier on RxList
quick, easy, pill identification
Use Pill Finder
Find a Local Pharmacy
including 24 hour, pharmacies
Find it Now
FACEBOOK
TWITTER
EMAIL
PRINT
Indications
Dosage
How Supplied
INDICATIONS
Treatment And Prevention Of Postmenopausal Osteoporosis
BONIVA is indicated for the treatment and prevention of osteoporosis in postmenopausal women. BONIVA increases bone mineral density (BMD) and reduces the incidence of vertebral fractures.
Important Limitations Of Use
The optimal duration of use has not been determined. The safety and effectiveness of BONIVA for the treatment of osteoporosis are based on clinical data of three years duration. All patients on bisphosphonate therapy should have the need for continued therapy re-evaluated on a periodic basis. Patients at low-risk for fracture should be considered for drug discontinuation after 3 to 5 years of use. Patients who discontinue therapy should have their risk for fracture re-evaluated periodically.
Last reviewed on RxList: 1/5/2015
• Osteopenia
• Osteoporosis and pain
• Bone density scans
Featured Centers
• 10 secrets to whiter teeth
• Does birth order determine who you are?
• Bad bugs in the UK: See pictures
ADVERTISEMENT
Move with us!
try our quote calculator >>
GLASSWELLS
TRUSTED REMOVALS & STORAGE
SINCE 1946
ADVERTISEMENT
Popular Collections

Figure 2: An example of the indications of 'Boniva' from `http://www.rxlist.com` website.

Table 2: Examples of medical concepts associated to 'osteoporosis' defined in our associated rules database.

Medical Concept	Related Concept
Dowager's hump	Osteoporosis
DEXA	Osteoporosis
Prolia	Osteoporosis
Boniva	Osteoporosis

heart disease' is a particular type of 'heart disease' (i.e. *'Coronary heart disease'→'heart disease'* – A patient suffering from 'Coronary heart disease' can be considered as having 'heart disease' condition).

For the free-text-based resources (e.g. `http://www.rxlist.com`), we use MetaMap to identify concepts related to the four aforementioned medical aspects from the free-text documents (we use the task-specific representation approach [17] described in Section 2 to identify medical concepts from sentences in the free-text documents), and then assume the existence of relationships between the identified concepts. For example, from a drug indication in the `http://www.rxlist.com` website, which states that *"Boniva (ibandronate sodium) is indicated for the treatment and prevention of osteoporosis in postmenopausal women"*, as shown in Figure 2, MetaMap can identify the concepts *'Boniva'* (treatment) and *'osteoporosis'* (diagnosis). Assuming the relationships between medical concepts found in the drug description, we surmise that there is an association between the two concepts (i.e. *'Boniva'→'osteoporosis'* and *'osteoporosis'→'Boniva'*). Next, the extracted association rules are stored in a database. For instance, as shown in Table 2, the rules associated to the concept *'osteoporosis'* are *'Dowager's hump'→'osteoporosis'*, *'DEXA'→'osteoporosis'*, *'Prolia'→'osteoporosis'*, and *'Boniva'→'osteoporosis'*. These association rules provide new evidence to infer the possible relevance of medical records, which are then aggregated to estimate the relevance of their associated patients. For instance, we can infer that patients taking the *'Boniva'* medicine suffer from *'osteoporosis'*, since *'Boniva'* is a treatment for *'osteoporosis'*.

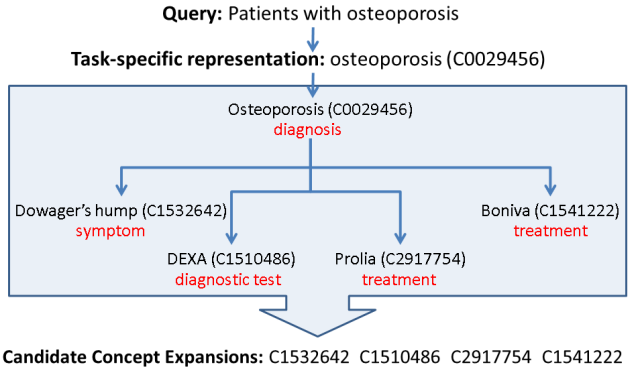

Figure 3: An example of identifying candidate concept expansions using our technique on query *"patients with osteoporosis"*.

3.2. A Bayesian-based Technique for Inferring Conceptual Relationships

We argue that the relationships between concepts that are related to the medical decision process could be leveraged to deal with the complexity of medical terminology. For example, if we have evidence that a patient is taking the *'olmesartan'* medicine (treatment), we can infer that the patient suffers from *'hypertension'* (diagnosis), since *'olmesartan'* is a treatment for *'hypertension'*. Therefore, using external domain-specific resources, we propose to reformulate the queries by using the association rules extracted from the medical resources (described in Section 3.1).

Specifically, we first retrieve a set of candidate concept expansions (denoted $inferred(q)$) corresponding to the query concepts from the extracted association rules. Then, to prevent excessively general candidate concepts being added to the query, we estimate the association of a query concept and each candidate concept expansion using a Bayesian probabilistic score computed based on the occurrences of concepts in the association rules (both derived from ontologies and free-text resources). The higher the probability, the stronger the relationship between the two concepts. Indeed, the Bayesian probabilistic score of the association between query concept t and its corresponding concept t' is estimated as follows:

$$w_a(t,t') = p(t'|t) = \frac{p(t' \cap t)}{p(t)} \tag{1}$$

where $p(t' \cap t)$ is the maximum likelihood that the concept t' co-occurs with the query concept t within all the extracted association rules, and $p(t)$ is the maximum likelihood that the concept t is contained in an association rule. This maximum likelihood is calculated based on the frequency of the occurrences of each concept in the whole association rules database.

Figure 3 shows an example of how our technique identifies candidate concept expansions for the query *"patients with osteoporosis"*. In particular, we first obtain the concept *'osteoporosis'* from the query using the task-specific representation approach. Then, the concept *'osteoporosis'* is used to retrieve related concepts from the database of association rules previously extracted from the medical resources (see Section 3.1). As shown in Figure 3, the retrieved candidate concept expansions include *'Dowager's hump'*, *'DEXA'*, *'Prolia'* and *'Boniva'*, which are the symptom, diagnostic test, and treatments associated to the original query concept (i.e. *'osteoporosis'*). However, as previously discussed in Section 2, QE is effective when a number of

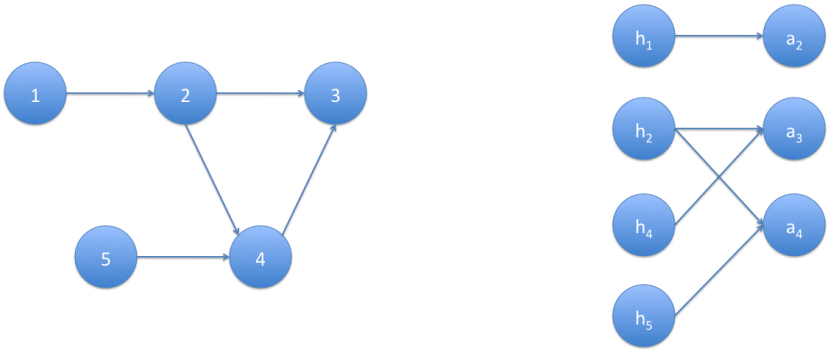

(a) A collection D of documents *1, 2, 3, 4, 5*

(b) A bipartite graph G

Figure 4: Transforming (a) the collection D into (b) a bipartite graph G.

highly informative terms that are associated to the query are added to the query [13], we select a top-k candidate concept expansions, which are ranked based on the score computed using Equation (1), to expand the original query.

3.3. A Stochastic Analysis for Inferring Conceptual Relationships

Next, we discuss our stochastic analysis technique for inferring relationships between medical concepts in the association rule database built in Section 3.1. Indeed, we propose to measure the strength of the relationships between the medical concepts in the association rule database using a stochastic analysis of some random walk behaviour through the association rules.

A stochastic analysis of random-walk behaviour has been studied before in IR (e.g. [31, 37, 38]). For example, [31]'s PageRank used a stochastic analysis of random walks on the hyperlink of the entire Web to determine the importance of each web page. [37] introduced a stochastic approach for link-structure analysis (SALSA), which performed random walks on a sub-graph derived from the link-structure of an initial set of retrieved results. In particular, consider a collection D of documents (e.g. *1, 2, 3, 4 and 5*) and directed relationships between the documents (see Figure 4(a)). For example, there is a relationship from document 1 to document 2. The approach first converts the collection D to a bipartite graph G, of which the two parts are considered *hubs* and *authorities*. Hubs are documents (i.e. nodes) that link to other documents, while authorities are documents that are linked to by other documents. The SALSA algorithm performs random walks to uncover the relationships between these documents.

Different from existing work, we propose to apply the SALSA algorithm to uncover medical concepts that are related to a given query and identify the strength of their relationships, using the link structure within our association rule database. Later, we discuss in Section 3.4 how these medical concepts are used to enrich the query.

Specifically, we adapt the SALSA algorithm to perform the analysis of the links between the concepts in the association rules. In this work, the nodes (e.g. *1, 2, 3, 4 and 5* in Figure 4(a)) are the medical concepts in the association rule database, while an edge is a directed relationship between two concepts (i.e. nodes). Consider a bipartite graph G where the two parts are hubs

Algorithm 1 The Adapted SALSA Algorithm

1: Initialise $a(s) \leftarrow 1$, $h(s) \leftarrow 1$ for all concepts $s \in C \cup C_a$, where C_a is a set of medical concepts in the association rule database that link to or are linked to by a concept in C.
2: Repeat the following steps until convergence:
3: Update the authority score of each concept s:
4: $a(s) \leftarrow \sum_{\{x|x \text{ points to } s\}} h(x)$
5: Update the hub score of each concept s:
6: $h(s) \leftarrow \sum_{\{x|s \text{ points to } x\}} a(x)$
7: Normalise the authority scores and the hub scores.

and authorities, a directed edge between hub h and authority a is indicated by a conceptual relationship between them (e.g. the edge between concept *1* and concept $2 - h_1 \rightarrow a_2$ in Figure 4(b)). A good hub points to many authorities, and a good authority is pointed by many hubs.

The authoritative medical concepts related to the query q should be linked to by many medical concepts in the sub-graph induced by a set of medical concepts C in the query q. A random walk on the sub-graph induced by C will visit those authoritative concepts with high probability. Hence, our adaptation of the SALSA algorithm [37] identifies medical concepts that are related to those medical concepts occurring in the query q, by using the iterative algorithm presented in Algorithm 1.

The medical concepts with high authority scores are the concepts highly related to the query q, as they can be inferred by many concepts related to the query. To incorporate these medical concepts while ranking patients based on the relevance of their medical records, we select top-k candidate concept expansions based on their authoritative scores computed using Algorithm 1.

3.4. Modelling Conceptual Relationships when Ranking Medical Records

We leverage both the local statistics from the top-ranked medical records and the association rules extracted from external resources to improve the representation of a given query. In particular, we apply a pseudo-relevance feedback technique to identify informative concepts from the top-ranked medical records retrieved using the original query. Indeed, when applying a pseudo-relevance feedback technique, concepts occurring in the top-ranked medical records are firstly weighted and ranked using a term weighting model. Then, the top-ranked concepts (i.e. the most informative concepts) are used to expand the original query. For example, for the query *"patients with vascular disease"*, which is represented as *"C0042373"* (*'vascular disease'*), a pseudo-relevance feedback QE technique could identify related concepts, such as *'C0190932'* (*'femoral-popliteal artery bypass graft'*) and *'C0014098'* (*'endarterectomy'*), which correspond to treatments for that disease. However, the related diagnostic procedures and symptoms, such as *'C0202896'* (*'carotid angiogram'*), might not be added to the query, if they do not appear in the top-ranked medical records. Hence, on other hand, we also expand the query with medical concepts inferred using the association rules. We identify candidate expansion concepts and their weights using either the Bayesian-based technique or the stochastic technique introduced in Section 3.2 or 3.3, respectively.

In particular, our approach estimates the relevance score of a medical record d towards the

query q as follows:

$$score(d,q) = \sum_{t'' \in q_e} qtw(t'') \cdot score(d,t'') + \lambda_r \cdot \sum_{t \in q} \sum_{t' \in inferred(q)} w_a(t,t') \cdot score(d,t') \quad (2)$$

where t'' and $qtw(t'')$ are a medical concept and its term weight in the expanded query q_e, which is reformulated using the occurrence statistics of medical concepts in the top-ranked medical records using any pseudo-relevance feedback model (e.g. Bo1 from the Divergence from Randomness (DFR) framework). $score()$ can be calculated using any term weighting model such as BM25, $inferred(q)$ returns a set of expanded concepts, which are related to the concepts in the original query q. We use either of the techniques previously proposed in Section 3.2 or 3.3 to extract the expanded concepts (i.e. $t' \in inferred(q)$) and their weights (i.e. $w_a()$) from the association rule database. In particular, the Bayesian-based technique, and the stochastic analysis technique uses Equation (1) and the authority score ($a()$) computed using Algorithm 1, respectively, to calculate $w_a()$. λ_r is a parameter to weight the importance of the relevance scores computed for the medical concepts derived using either of the two proposed techniques.

To learn the λ_r value for each query, we use the Gradient Boosted Regression Trees (GBRT) regression technique [39] (as implemented in the jforest package [40][8]), since it has shown to be effective for regression tasks (e.g. [41]). To learn an value of parameter λ_r for an unseen query, we use 12 features, which measure the predicted difficulty of the query. An effective feature should indicate the level of emphasis on the relevance scores computed for the concepts inferred using the association rules. Our intuition for using these features is that if the query is difficult, then it might be beneficial to take into account the relevance estimated from the inferred concepts. In particular, if using only the information from the original query is difficult for a search system, medical concepts inferred using the extracted association rules might bring novel evidence that can improve the representation of the query and hence enhance the retrieval performance. Table 3 lists the 12 query performance predictors computed on the original query, which are well-known for measuring the difficulty of a query. Specifically, the first set of features (Features 1-4), including the clarity score [42], SCQ [43], MaxSCQ [43] and NSCQ [43], consider the ambiguity of a query by measuring the coherence of the language used in each medical record. The more similar the query model is to the collection model, the better the retrieval performance would be expected. The next set of features (Features 5-8) measure the specificity of each query. Indeed, queries with explicit intents could result in a better performance than queries with general concepts. The features include Average Inverse Collection Term Frequency (AvICTF) [44], Average Inverse Document Frequency (AvIDF) [44], EnIDF [44], and the query scope (ω) [45]. Features 9-10 measure the distribution of informativeness among the query concepts (i.e. γ_1 and γ_2 [45]), as a query with informative concepts could attain an effective retrieval performance. Next, Feature 11, the Average of the Pointwise Mutual Information over all query term pairs (AvPMI) [44], focuses on the relationship between query concepts. The more co-occurrences among the query concepts, the better the chance that the relevant documents are being retrieved. Finally, Feature 12 is the number of non-stopword query concepts.

[8] `http://code.google.com/p/jforests`

Table 3: List of the features used for learning the emphasise on medical concepts derived from association rules.

ID	Feature
1	Clarity Score [42]
2	SCQ [43]
3	MAXCQ [43]
4	NSCQ [43]
5	AvICTF [44]
6	AvIDF [44]
7	EnIDF [44]
8	Query Scope (ω) [45]
9	γ_1 [45]
10	γ_2 [45]
11	AvPMI [44]
12	Query length [45]

4. Experimental Setup

4.1. Corpus/Queries/Measures

We evaluate our proposed approach in the context of the TREC 2011 [4] and 2012 [5] Medical Records track. In this track, the task is to identify relevant patient *visits* for a given query topic. Each visit contains all of the medical records associated with a patient's visit to a hospital. A visit is used to represent a *patient* as a unit of retrieval, due to privacy concerns [4]. The medical records collection consists of approximately 102k medical records, which can be mapped to 17,265 patient visits, from the University of Pittsburgh NLP Repository. We evaluate our proposed approach using the 34 and 47 topics from the TREC 2011 and 2012 Medical Records track, respectively.

We evaluate the retrieval effectiveness of our proposed approach, in terms of the TREC official measures, which are bpref measure [46] for TREC 2011, and infNDCG measure [47] for TREC 2012. In particular, bpref is the official measure of TREC 2011, since the absolute number of judged patient visits per topic is relatively small [4]. bpref is designed for evaluating environments with incomplete relevance data and penalises a system which ranks a judged non-relevant document above a judged relevant document [46]. On the other hand, infNDCG is the primary measure of the TREC 2012 Medical Records track [5], since the gold standard judgements are incomplete; hence a sampling approach is deployed to infer the NDCG performance. In addition, we measure statistically significant differences between the retrieval performance achieved by our approach and the baselines using the paired t-test at $p < 0.05$.

4.2. Medical Record & Patient Ranking

We index the medical records using the Terrier retrieval platform [48]. The BM25 weighting model [26] is used to rank medical records (e.g. $score_t()$ in Equation (2)). Then, to rank the patients, as explained in Section 2, we deploy the expCombSUM voting technique [24], which gives more importance to the highly relevant medical records while voting for the relevance of the patients. In particular, for a given ranking of medical records ($R(Q)$) with respect to query Q, each medical record is said to vote for the relevance of its associated patient. The number of medical records in $R(Q)$ to vote for the relevance of the patients is limited to 5,000, as suggested in [21].

Table 4: List of resources used for extracting the conceptual relationships related to the four aspects of the medical decision process.

Resource	Description of the Extracted Association Rules
DOID hierarchy	Hierarchical relationships between concepts within the same aspects e.g. a general disease and a specific type of the general disease
MeSH	Hierarchical relationships between concepts within the same aspects e.g. a general disease and a specific type of the general disease
MedDRA	Hierarchical relationships between concepts within the same aspects e.g. a general disease and a specific type of the general disease
DOID	Relationships between concepts across the aspects e.g. a disease and its symptoms
`http://www.rxlist.com`	Relationships between concepts across the aspects e.g. a medicine and the diseases that it can remedy
`http://www.webmd.com`	Relationships between concepts across the aspects e.g. a diagnostic test and the diseases that it can diagnosed
UMLS	Both hierarchical and across-aspects relationships e.g. a general disease and a specific type of the general disease

4.3. Association Rules

To create the association rule database (discussed in Section 3.1), we use the external resources listed in Table 4, which are representatives of both ontology-based and free-text-based medical resources. Table 5 shows the number of association rules between concepts extracted from the 7 domain-specific resources. The types of relationships that are extracted from each resource are described in Column 2 of Table 4. Note that there are some association rules that overlap between resources. In total, there are 11,373,014 extracted association rules in our database. Meanwhile, Table 6 shows the number of the extracted association rules with respect to the four aspects of the medical decision process. For example, the rules with the type 'diagnosis $\rightarrow$ symptom' are the association rules that can be used to infer the related symptoms that the medical records of the relevant patients are likely to contain if the query contains a particular diagnosis. Note that there are some association rules that are duplicated among different types of rules, since using the technique described in [17], some medical concepts can be categorised into several medical aspects. For example, medical concepts related to parts of body are considered to be associated with symptom, diagnostic test, diagnosis and treatment (see Table 1). We use these extracted association rules to extract concept expansions by using the Bayesian-based and stochastic techniques previously introduced in Sections 3.2 and 3.3, respectively. In addition, we follow [13] and use only *top 10* candidate concept expansions ($k = 10$) extracted using each technique.

For the local-statistic query expansion technique used in Equation (2) of Section 3.4, we deploy a parameter-free Bose-Einstein statistics-based (Bo1) model from the DFR framework to extract informative concepts from the top-ranked medical records as it has shown to be effective in previous works [21, 17]. DFR Bo1 measures the informativeness of a concept based on the divergence of the concept's distribution in the pseudo-relevant feedback documents (i.e. top-

Table 5: Number of association rules extracted from each domain-specific resource.

Resources	Association Types	# of Rules
DOID hierarchy	Specific-general	2,046
MeSH	Specific-general	53,915
MedDRA	Specific-general	86,109
DOID	Across aspects	9,664
`http://www.rxlist.com`	Across aspects	5,433
`http://www.webmd.com`	Across aspects	3,694
UMLS	Specific-general & Across aspects	11,212,153

Table 6: The types of the extracted association rules in terms of the four medical aspects.

Rule type	# of rules
diagnosis → diagnosis	2,900,788
diagnosis → diagnostic test	1,907,492
diagnosis → symptom	2,306,712
diagnosis → treatment	2,018,061
diagnostic test → diagnosis	1,905,521
diagnostic test → diagnostic test	2,726,341
diagnostic test → symptom	2,023,443
diagnostic test → treatment	2,695,757
symptom → diagnosis	2,303,486
symptom → diagnostic test	2,020,294
symptom → symptom	4,205,071
symptom → treatment	2,096,350
treatment → diagnosis	2,015,362
treatment → diagnostic test	2,698,453
treatment → symptom	2,097,231
treatment → treatment	6,597,028

ranked medical records) from its distribution in the entire collection. The more the divergence is measured, the more likely that the concept is informative and related to the query. Specifically, Bo1 from the Divergence from Randomness framework calculates the weight $w(t)$ of a concept t in the pseudo-relevant documents, as follows [13]:

$$w(t) = tf_x \cdot log_2 \frac{1+P_n}{P_n} + log_2(1+P_n) \tag{3}$$

where tf_x is the frequency of concept t in the set of pseudo-relevant documents x. $P_n = \frac{TF}{N}$, where TF is the term frequency of t in all documents in the collection, and N is the number of documents in the collection. In this work, we follow [13] and set the number of pseudo-relevant documents and expanded concepts to 3 and 10, respectively.

Once we identify the expanded concepts, a parameter-free function is used to calculate the

query term weight $qtw(t)$ of each concept t in the expanded query, as follows [13]:

$$qtw(t) = \frac{qtf}{qtf_{max}} + \frac{w(t)}{\lim_{TF \to tfw} w(t)} \tag{4}$$
$$= TF_{max} \log_2 \frac{1 + P_{n,max}}{P_{n,max}} + \log_2(1 + P_{n,max})$$

where qtf is the frequency of a concept t in the query, and qtf_{max} is the maximum of qtf. $\lim_{TF \to tfw} w(t)$ is the upper bound of $w(t)$, TF_{max} is the frequency TF of the concept with the maximum $w(t)$ in the pseudo-relevant documents. $P_{n,max} = \frac{TF_{max}}{N}$. Note that if an original query concept t does not appear in the most informative concept list extracted from the pseudo-relevant documents, its query term weight $qtw(t)$ remains equal to the original one.

As both the Bayesian-based and the stochastic analysis-based techniques require the λ_r parameter in Equation (2) to be properly set, we apply two different training regimes. Firstly, we apply a five-fold cross-validation (5-fold) within each set of queries, since the target measure of TREC 2011 and 2012 queries are different. Secondly, we use the TREC 2012 queries as training data, when testing with the TREC 2011 queries, and vice versa. We refer to this setting as *cross-collection validation* (x-collection). We optimise the TREC primary measure (i.e. bpref and infNDCG measures for TREC 2011 and 2012, respectively) when learning the λ_r parameter.

4.4. Baselines

We compare the retrieval effectiveness of the proposed approach to infer relationships between medical concepts with the following baselines:

- The task-specific representation baseline [17] (see Section 2), which does not apply any QE technique.
- The Bo1 baseline: applying the Bo1 query expansion technique with the task-specific representation approach.
- The Bo1 & Semantic QE baseline: using both the Bo1 query expansion technique and a semantic query expansion technique that adds all of the related medical concepts into the queries (as used in [36]). Note that this baseline uses the same set of external resources as the proposed approach.

5. Experimental Results

In Sections 5.1 and 5.2, we evaluate the proposed approach when using Bayes' theorem and a stochastic analysis to derive associated medical concepts t'' and their weights $qtw(t'')$ in Equation (2), respectively. In particular, we refer the two variants as the Bayesian-based approach and the stochastic approach.

5.1. Experiment with our Bayesian-based Approach for Inferring Conceptual Relationships

We first evaluate the effectiveness of the Bayesian-based technique for inferring the relationships between medical concepts. To have a fair train/test setting, we use the *cross-collection validation* and *5-fold cross validation* as described in Section 4 when setting λ_r within the Bayesian-based approach. Furthermore, to see how the parameter setting impacts on the retrieval performance and the potential effectiveness of our approach, the performance achieved when using

Table 7: Retrieval performance of our Bayesian-based approache for inferring relationships between medical concepts in comparison with the three baselines on TREC 2011 and 2012 Medical Records track's queries. Statistical significance (paired t-test) at $p < 0.05$ over the task-specific representation and the Bo1 & Semantic QE baselines are denoted ^{t} and ^{s}, respectively.

Approach	bpref (TREC 2011)	infNDCG (TREC 2012)
Task-specific representation	0.5243	0.4880
+ Bo1	0.5403	0.5129
+ Bo1 & Semantic QE	0.3013	0.2228
+ Our Bayesian-based approach (5-fold)	**0.5474**^{s}	**0.5133**^{t,s}
+ Our Bayesian-based approach (x-collection)	0.5464^s	0.5125^{t,s}
+ Our Bayesian-based approach (oracle)	0.5646^{t,s}	0.5393^{t,s}

the best λ_r for each retrieval measure on each test collection (i.e. best possible setting) is also reported, denoted *oracle*.

Table 7 compares the retrieval performances of the Bayesian-based approach with the three baselines on TREC 2011 and 2012 Medical Records track test collection. From Table 7, we observe that the Bo1 baseline outperforms the task-specific representation baseline (i.e. no QE) for both TREC 2011 and 2012. Specifically, the retrieval performances improve from bpref 0.5243 to 0.5403 and from infNDCG 0.4880 to 0.5129 for TREC 2011 and 2012, respectively. Meanwhile, the Bo1 & Semantic QE baseline decreases the retrieval performance.

In contrast, the Bayesian-based approach, using either a 5-fold or an x-collection training regime, outperforms the task-specific representation baseline (both TREC 2011 and 2012). Specifically, the Bayesian-based approach with the 5-fold regime (5-fold) outperforms all of the three used baselines. The Bayesian-based approach (5-fold) achieves a bpref of 0.5474 and an infNDCG of 0.5133. Importantly, our approach (5-fold) significantly (paired t-test, $p < 0.05$) outperforms the task-specific representation baseline in terms of infNDCG for TREC 2012 (0.5133 vs 0.4880).

When comparing the Bayesian-based approach with the Bo1 & Semantic QE baseline, which also uses the same set of resources as the Bayesian-based approach, we observe that the Bayesian-based approach, either using the 5-fold or x-collection training regime, significantly ($p < 0.05$) outperforms this baseline for both the TREC 2011 and 2012 queries. However, the Bayesian-based approach could not significantly outperform the Bo1 baseline.

Additionally, as expected, we find that with a proper setting of the parameter λ_r (i.e. the λ_r that results in the best retrieval performance for each query), the Bayesian-based approach (oracle) can achieve a better retrieval performance. In particular, the bpref retrieval performance is increased to 0.5646, while the infNDCG retrieval performance is improved to 0.5393 (+7.69% and +10.51% over the task-specific representation baseline, respectively). This shows the potential of the Bayesian-based approach if the λ_r value could be effectively set.

Later, in Section 6, we further analyse the Bayesian-based approach, in comparison with the task-specific representation and the Bo1 baselines.

5.2. Experiments with a Stochastic Approach for Inferring Conceptual Relationships

Next, we discuss the evaluation of the stochastic approach for inferring the conceptual relationships to improve the representation of a given query. We compare the retrieval effectiveness of our approach with the same baselines as in Section 5.1. Table 8 shows the retrieval performances in terms of bpref and infNDCG for TREC 2011 and 2012, respectively.

Table 8: Retrieval performance of our stochastic approach for inferring relationships between medical concepts in comparison with the three baselines on TREC 2011 and 2012 Medical Records track's queries. Statistical significance (paired t-test) at $p < 0.05$ over the task-specific representation and the Bo1 & Semantic QE baselines are denoted ^{t} and ^{s}, respectively.

Approach	bpref (TREC 2011)	infNDCG (TREC 2012)
Task-specific representation	0.5243	0.4880
+ Bo1	**0.5403**	**0.5129**
+ Bo1 & Semantic QE	0.3013	0.2228
+ Our stochastic approach (5-fold)	0.5355^s	0.4846^s
+ Our stochastic approach (x-collection)	0.5376^s	0.4761^s
+ Our stochastic approach (oracle)	$0.5586^{t,s}$	$0.5311^{t,s}$

From Table 8, we observe that the stochastic approach significantly ($p < 0.05$) outperforms the Bo1 & Semantic QE baseline for both TREC 2011 and 2012 queries. In addition, the stochastic approach outperforms the task-specific representation baseline for the TREC 2011 queries. Specifically, our approach with the 5-fold and the x-collection training regimes achieves bpref 0.5355 and 0.5376, respectively, in comparison with bpref 0.5243 for the task-specific representation baseline. However, for the TREC 2012 queries, our approach (either with the 5-fold or the x-collection training regime) could not improve the retrieval performance over the task-specific representation baseline. Meanwhile, we observe that the Bo1 baseline performs better than the stochastic approach. This is partly due to the limited number of available queries for training the λ_r parameter (in Equation (2)).

When considering the oracle setting with λ_r set to the most effective value for each query, we observe that the stochastic approach could further improve the retrieval performance significantly. This shows the potential of the stochastic approach. However, we leave for future work the study of a more effective training of λ_r.

Next, when comparing the two proposed approaches (see Tables 7 and 8), we observe that the Bayesian-based approach is more effective than the stochastic approach for both TREC 2011 and TREC 2012. For example, with the 5-fold training regime, the Bayesian-based approach achieves infNDCG 0.5133, while the infNDCG performance of the stochastic approach is 0.4846.

6. Analysis and Discussion

This section further discusses the performances of the Bayesian-based and the stochastic approaches for inferring the relationships between medical concepts. Specifically, Section 6.1 discusses which types of queries are likely to benefit from these proposed approaches, while Sections 6.2 and 6.3 compare the performances of the two proposed approaches with the task-specific representation baseline and the Bo1 baseline, respectively, on a per-query basis.

6.1. Failure Analysis

In this section, we analyse the performances of both the Bayesian-based and the stochastic approaches in comparison with the task-specific representation baseline to identify which types of queries that are likely to benefit from each of the two approaches. Tables 9 and 10 show the numbers of queries benefited or harmed by the Bayesian-based and the stochastic approaches, respectively, which are grouped based on the types of the medical concepts in the queries.

Table 9: Analysis of our Bayesian-based approach w.r.t. the aspects of medical concepts found in the queries. The numbers between the parentheses indicate the number of queries impacted (benefited/harmed) compared to the total number of queries.

Aspects of medical concepts in the query	Benefited	Harmed
Symptom	52.9% (37/70)	42.9%(30/70)
Diagnostic test	57.7% (30/52)	38.5%(20/52)
Diagnosis	55.9% (38/68)	38.2% (26/68)
Treatment	55.8% (29/52)	40.4% (21/52)
All 4	50% (16/32)	46.9%(15/32)

Table 10: Analysis of our stochastic approach w.r.t. the aspects of medical concepts found the queries. The numbers between the parentheses indicate the number of queries impacted (benefited/harmed) compared to the total number of queries.

Aspects of medical concepts in the query	Benefited	Harmed
Symptom	52.9% (37/70)	41.4%(29/70)
Diagnostic test	55.8% (29/52)	40.4%(21/52)
Diagnosis	51.5% (35/68)	41.2% (28/68)
Treatment	53.8% (28/52)	38.5% (20/52)
All 4	50% (16/32)	43.8%(14/32)

From Tables 9 and 10, we observe that both the Bayesian-based approach and the stochastic approach are more likely to be effective when the query contains medical concepts related to diagnostic tests. In particular, the Bayesian-based approach and the stochastic approach improve the retrieval performance for 57.7% and 55.8% of the queries containing medical concepts related to diagnostic tests, respectively. Meanwhile, we observe that the queries that contain medical concepts related to all of the four aspects are less likely to benefit from our two approaches.

6.2. Comparison with the Task-Specific Representation Baseline

Next, we compare the retrieval performance of the Bayesian-based and the stochastic approaches with the task-specific representation baseline on a per-query basis. We choose to discuss only the 5-fold setting as it is the most effective.

We first discuss the performance of the Bayesian-based approach. From Figure 5, we observe that in general the Bayesian-based approach performs better than the task-specific representation baseline. In particular, for the TREC 2011 queries, the Bayesian-based approach outperforms the task-specific representation baseline for 22 out of 34 queries, while it performs worse than the baseline for 10 queries. We observe that the Bayesian-based approach is likely to be more effective for difficult queries. For instance, the Bayesian-based approach could improve the retrieval performance of the queries with a bpref retrieval performance less than 0.25 for 5 out of 6 queries (e.g. queries# 108, 111, 121, and 125). Meanwhile, for the TREC 2012 queries, the Bayesian-based approach outperforms the task-specific representation baseline for 23 out of 47 queries, while performing worse than the baseline for 22 queries.

In Figure 6, we show the per-query retrieval effectiveness of the stochastic approach in comparison with the task-specific representation baseline. We observe that the stochastic approach outperforms this baseline for the majority of the queries. Indeed, our approach is more effective for 18 out of 34 queries and for 24 out of 47 queries from TREC 2011 and 2012, respectively. Similar to the Bayesian-based approach, we also observe that the stochastic approach tends to be

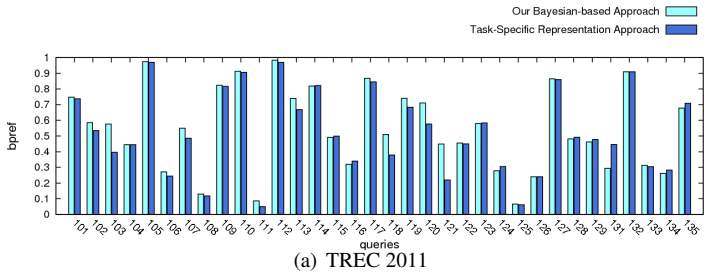

(a) TREC 2011

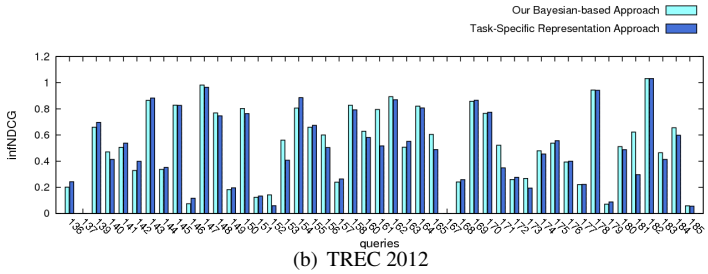

(b) TREC 2012

Figure 5: The retrieval performances of our Bayesian-based approach for inferring the conceptual relationships and the task-specific representation baseline on a per-query basis, evaluated using the TREC 2011 and 2012 Medical Records Track.

effective for difficult queries. Specifically, for the queries where the task-specific representation baseline obtains a retrieval performance less than 0.25 (bpref for TREC 2011 and infNDCG for TREC 2012), the stochastic approach outperforms the task-specific representation baseline for 7 out of 11 queries for TREC 2011 and for 5 out of 6 for TREC 2012.

6.3. Comparison with the Bo1 Baseline

This section discusses the retrieval performances of the Bayesian-based and the stochastic approaches in comparison with the Bo1 baseline, which is a stronger baseline. We compare the retrieval performances on a per-query basis.

We first discuss the retrieval effectiveness of the Bayesian-based approach. From Figure 7, we observe that the Bayesian-based approach performs comparable to the Bo1 baseline. Indeed, our approach outperforms the Bo1 baseline for 16 out of 34 queries and for 21 out of 47 queries for TREC 2011 and TREC 2012, respectively. Meanwhile, the Bo1 baseline performs better than our approach for 12 and 20 queries for TREC 2011 and 2012.

Next, we discuss the retrieval performance of the proposed stochastic approach. From Figure 8, we observe that the Bo1 baseline benefits more queries than the stochastic approach. Specifically, the Bo1 baseline outperforms the stochastic approach for 16 (resp. 23) out of 34 (resp. 47) queries for TREC 2011 (resp. TREC 2012), while the stochastic approach performs better for 11 and 20 queries for TREC 2011 and 2012, respectively. However, this is expected, as the Bo1 baseline achieves a better retrieval performance than the stochastic approach, as shown in Table 8. Nevertheless, when considering the performance on *the difficult queries* (i.e. queries that obtain retrieval performance < 0.25 when query expansion is not applied), we observe that the stochastic approach performs better than the Bo1 baseline for 8 out of 16 queries across the TREC 2011 and 2012, while it performs worse than the Bo1 baseline for only 3 queries.

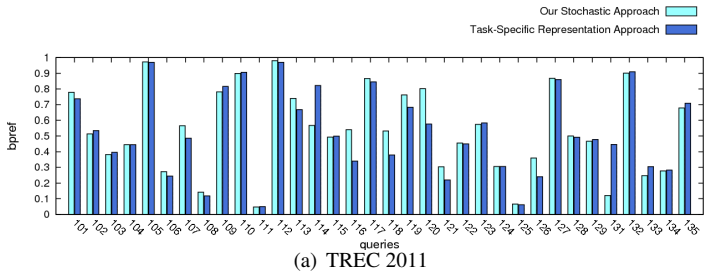

(a) TREC 2011

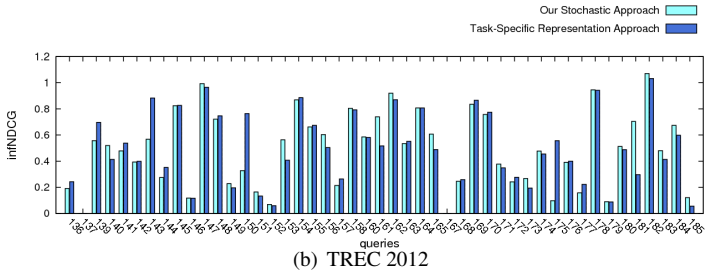

(b) TREC 2012

Figure 6: The retrieval performances of our stochastic approach for inferring the conceptual relationships and the task-specific representation baseline on a per-query basis, evaluated using the TREC 2011 and 2012 Medical Records Track.

From the per-query performance comparison in Sections 6.2 and 6.3, we observe that both the Bayesian-based and the stochastic approaches are more likely to benefit the difficult queries (i.e. the queries that obtain retrieval performance < 0.25 when do not apply query expansion).

7. Conclusions

We have proposed a novel approach that uncovers implicit knowledge in medical records and queries. In particular, the proposed approach to infer the relationships between medical conditions by using both external resources (e.g. ontologies and health-related websites) and the local statistics of top-ranked medical records to reformulate the queries. In particular, we extracted association rules of relationships between medical concepts from external resources. Then, to infer the relationships of medical concepts from the association rules, we introduced two techniques based on Bayes' theorem and a stochastic analysis, respectively. The medical concepts inferred using the association rules are combined with the medical concepts derived using a pseudo-relevance feedback technique in order to improve the representation of the query. From the experimental results, we observed that with a fair setting the Bayesian-based approach improved the retrieval performance by up to 10.51% over the baseline that does not apply a QE technique. Meanwhile, the stochastic approach was likely to perform effectively for difficult queries (i.e. the queries where the retrieval performance of the original queries are less than 0.25 in term of bpref and infNDCG for TREC 2011 and 2012, respectively). In addition, we observed that both the Bayesian-based and the stochastic approaches were less likely to be effective for the queries that contain medical concepts related to all of the four aspects of the medical decision process. For future work, we aim to investigate a more effective technique for extracting association rules from different external resources and compare which resources are more useful for

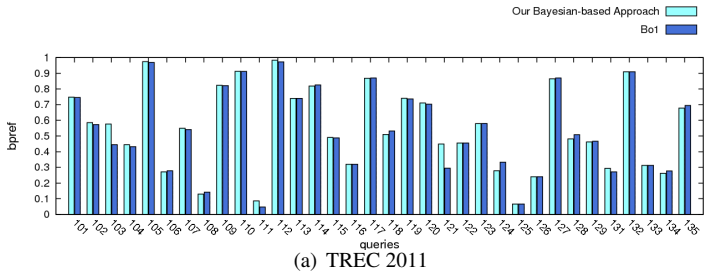

(a) TREC 2011

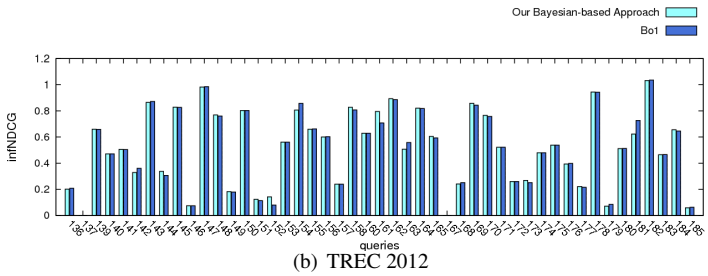

(b) TREC 2012

Figure 7: The retrieval performances of our Bayesian-based approach for inferring the conceptual relationships and the Bo1 query baseline on a per-query basis, evaluated using the TREC 2011 and 2012 Medical Records Track.

the patient search task.

References

[1] I. A. Kotsiopoulos, J. A. Keane, M. Turner, P. J. Layzell, and F. Zhu. Ibhis: Integration broker for heterogeneous information sources. In *Proceedings of the 27th Annual International Computer Software and Applications Conference*, 2003.

[2] E. Tambouris and C. Makropoulos. Hin6/427: Co-operative health information network in europe: The greek experience. *Journal of Medical Internet Research*, 2(2), Sep 1999.

[3] N. Limsopatham, C. Macdonald, and I. Ounis. Inferring conceptual relationships to improve medical records search. In *Proceedings of the 10th Open research Areas in Information Retrieval*, Lisbon, Portugal, 2013.

[4] E. Voorhees and R Tong. Overview of the TREC 2011 medical records track. In *Proceedings of the 20th Text REtrieval Conference*, Gaithersburg, MD, USA, 2011.

[5] E. Voorhees and W Hersh. Overview of the TREC 2012 medical records track. In *Proceedings of the 21st Text REtrieval Conference*, Gaithersburg, MD, USA, 2012.

[6] Michael Krauthammer and Goran Nenadic. Term identification in the biomedical literature. *J. of Biomedical Informatics*, 37(6):512–526, December 2004.

[7] N. Limsopatham, R. L. T. Santos, C. Macdonald, and I. Ounis. Disambiguating biomedical acronyms using emim. In *Proceedings of the 34th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Beijing, China, 2011.

[8] D. Trieschnigg, D. Hiemstra, F. de Jong, and W. Kraaij. A cross-lingual framework for monolingual biomedical information retrieval. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*, New York, NY, USA, 2010.

[9] W. Hersh, D. Hickam, R. Haynes, and K. McKibbon. A performance and failure analysis of saphire with a medline test collection. *Journal of the American Medical Informatics Association*, 1(1), 1994.

[10] P. Srinivasan. Optimal document-indexing vocabulary for medline. *Information Processing and Management*, 32(5), September 1996.

[11] A. R. Aronson and T. C. Rindflesch. Query expansion using the UMLS Metathesaurus. *Proceedings of the American Medical Informatics Association (AMIA) Symposium*, 1997.

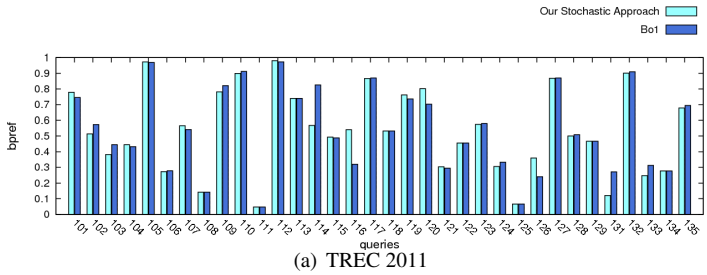

(a) TREC 2011

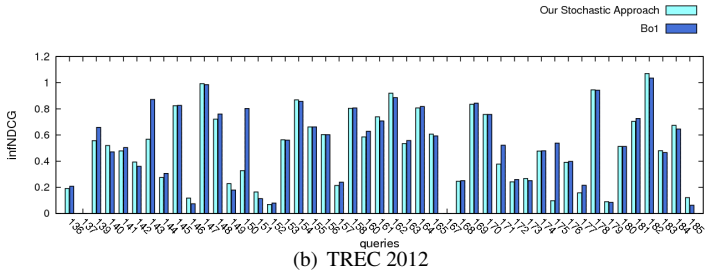

(b) TREC 2012

Figure 8: The retrieval performances of our stochastic approach for inferring the conceptual relationships and the Bo1 baseline on a per-query basis, evaluated using the TREC 2011 and 2012 Medical Records Track.

[12] Padmini Srinivasan. Query expansion and medline. *Inf. Process. Manage.*, 32(4):431–443, July 1996.

[13] G. Amati. *Probabilistic Models for Information Retrieval based on Divergence from Randomness*. PhD thesis, University of Glasgow, 2003.

[14] N. Limsopatham, C. Macdonald, R. McCreadie, and I. Ounis. Exploiting term dependence while handling negation in medical search. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval*, New York, NY, USA, 2012.

[15] Nut Limsopatham, Craig Macdonald, and Iadh Ounis. Modelling the usefulness of document collections for query expansion in patient search. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, CIKM '15, pages 1739–1742, New York, NY, USA, 2015. ACM.

[16] E. Silfen. Documentation and coding of ed patient encounters: an evaluation of the accuracy of an electronic medical record. *The American Journal of Emergency Medicine*, 24(6), 2006.

[17] N. Limsopatham, C. Macdonald, and I. Ounis. A task-specific query and document representation for medical records search. In *Proceedings of the 35th European Conference on Information Retrieval*, Moscow, Russia, 2013.

[18] W. Hersh. *Information Retrieval: A Health and Biomedical Perspective (Health Informatics)*. Springer, 3rd edition, November 2008.

[19] W. Hersh. Health care information technology: Progress and barriers. *JAMA: The Journal of the American Medical Association*, 292(18), 2004.

[20] H. Gurulingappa, B. Uller, M. Hofmann-Apitius, and J. Fluck. A semantic platform for information retrieval from e-health records. In *Proceedings of the 20th Text REtrieval Conference*, Gaithersburg, MD, USA, 2011.

[21] N. Limsopatham, C. Macdonald, I. Ounis, G. McDonald, and M. Bouamrane. University of glasgow at medical records track 2011: Experiments with terrier. In *Proceedings of the 20th Text REtrieval Conference*, Gaithersburg, MD, USA, 2011.

[22] D. Zhu and B. Carterette. Combining multi-level evidence for medical record retrieval. In *Proceedings of the 2012 International Workshop on Smart Health and Wellbeing*, New York, NY, USA, 2012.

[23] K. Balog, P. Thomas, N. Craswell, I. Soboroff, , and P. Bailey. Overview of the trec 2008 enterprise track. In *Proceedings of the 17th Text REtrieval Conference*, Gaithersburg, MD, USA, 2008.

[24] C. Macdonald and I. Ounis. Voting for candidates: adapting data fusion techniques for an expert search task. In *Proceedings of the 15th ACM international conference on Information and knowledge management*, New York, NY, USA, 2006.

[25] K. Balog, L. Azzopardi, and M. de Rijke. Formal models for expert finding in enterprise corpora. In *Proceedings*

of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Seattle, Washington, USA, 2006.

[26] S. E. Robertson and S. Walker. Some simple effective approximations to the 2-poisson model for probabilistic weighted retrieval. In *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, New York, NY, USA, 1994.

[27] G. Amati, E. Ambrosi, M. Bianchi, C. Gaibisso, and G. Gambosi. Fub, iasi-cnr and university of tor vergata at trec 2007 blog track. In *Proceedings of the 16th Text REtrieval Conference*, Gaithersburg, MD, USA, 2007.

[28] A. R. Aronson. Exploiting a large thesaurus for information retrieval. In *Proceedings of RIAO 1994*, 1994.

[29] A. R. Aronson and F. M. Lang. An overview of metamap: historical perspective and recent advances. *Journal of the American Medical Informatics Association*, 17(3), 2010.

[30] B. Koopman, G. Zuccon, P. Bruza, L. Sitbon, and M. Lawley. Graph-based concept weighting for medical information retrieval. In *Proceedings of the Seventeenth Australasian Document Computing Symposium*, New York, NY, USA, 2012.

[31] S. Brin and L. Page. The anatomy of a large-scale hypertextual web search engine. In *Proceedings of the Seventh International Conference on World Wide Web 7*, Amsterdam, The Netherlands, The Netherlands, 1998.

[32] B. Koopman and G. Zuccon. Understanding negation and family history to improve clinical information retrieval. In *Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval*, New York, NY, USA, 2014.

[33] N. Limsopatham, C. Macdonald, and I. Ounis. Modelling relevance towards multiple inclusion criteria when ranking patients. In *Proceedings of the 23rd ACM International Conference on Information and Knowledge Management*, Shanghai, China, 2014.

[34] David Martinez, Arantxa Otegi, Aitor Soroa, and Eneko Agirre. Improving search over electronic health records using umls-based query expansion through random walks. *Journal of Biomedical Informatics*, 51:100 – 106, 2014.

[35] Guido Zuccon, Bevan Koopman, Anthony Nguyen, Deanne Vickers, and Luke Butt. Exploiting medical hierarchies for concept-based information retrieval. In *Proceedings of the Seventeenth Australasian Document Computing Symposium*, pages 111–114. ACM, 2012.

[36] B. King, L. Wang, I. Provalov, and J. Zhou. Cengage learning at trec 2011 medical track. In *Proceedings of the 20th Text REtrieval Conference*, Gaithersburg, MD, USA, 2011.

[37] R. Lempel and S. Moran. Salsa: The stochastic approach for link-structure analysis. *ACM Transactions on Information Systems*, 19(2), April 2001.

[38] R. Blanco and C. Lioma. Graph-based term weighting for information retrieval. *Information Retrieval*, 15(1), February 2012.

[39] S. Tyree, K. Q. Weinberger, K. Agrawal, and J. Paykin. Parallel boosted regression trees for web search ranking. In *Proceedings of the 20th International Conference on World Wide Web*, New York, NY, USA, 2011.

[40] Y. Ganjisaffar, R. Caruana, and C. V. Lopes. Bagging gradient-boosted trees for high precision, low variance ranking models. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*, New York, NY, USA, 2011.

[41] N. Limsopatham, C. Macdonald, R. McCreadie, and I. Ounis. Learning to combine representations for medical records search. In *Proceedings of the 36th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Dublin, Ireland, 2013.

[42] S. Cronen-Townsend, Y. Zhou, and W. B. Croft. Predicting query performance. In *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, New York, NY, USA, 2002.

[43] Y. Zhao, F. Scholer, and Y. Tsegay. Effective pre-retrieval query performance prediction using similarity and variability evidence. In *Proceedings of the IR Research, 30th European Conference on Advances in Information Retrieval*, Berlin, Heidelberg, 2008.

[44] D. Carmel and E. Yom-Tov. *Estimating the Query Difficulty for Information Retrieval*. Synthesis Lectures on Information Concepts, Retrieval, and Services. Morgan & Claypool Publishers, 2010.

[45] B. He and I. Ounis. Query performance prediction. *Information Systems*, 31(7), November 2006.

[46] C. Buckley and E. M. Voorhees. Retrieval evaluation with incomplete information. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, New York, NY, USA, 2004.

[47] E. Yilmaz, E. Kanoulas, and J. A. Aslam. A simple and efficient sampling method for estimating ap and ndcg. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, New York, NY, USA, 2008.

[48] I. Ounis, G. Amati, V. Plachouras, B. He, C. Macdonald, and C. Lioma. Terrier: A High Performance and Scalable Information Retrieval Platform. In *Proceedings of ACM SIGIR'06 Workshop on Open Source Information Retrieval (OSIR 2006)*, 2006.